



ZPW
AF

PATENT

I hereby certify that on the date specified below, this correspondence is being deposited with the United States Postal Service as first-class mail in an envelope addressed to Mail Stop Appeal Brief, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

05.23-05

Date


Joanne Bourguignon

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventors: Robert Alan Cochran
Serial No. 09/678,168
Filed: October 02, 2000
For: Method and System for Throttling I/O Request Servicing on Behalf of an I/O Request Generator to Prevent Monopolization of a Storage Device by the I/O Request Generator

Examiner: Akiba K.R. Boyce
Group Art Unit: 3623
Docket No. 10992806-1
Date: May 23, 2005

Mail Stop Appeal Brief
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REQUEST FOR REINSTATEMENT OF THE APPEAL UNDER 37 CFR § 1.193(2)(ii)

Sir:

In response to the Appeal Brief filed November 30, 2004 and the subsequent Office Action dated February 22, 2005, Applicant respectfully requests the reinstatement of the appeal. Accompanying this Request is a Supplemental Appeal Brief in triplicate.

Applicant believes that no fee is required. However, at any time during the pendency of this application, please charge any fees required or credit any overpayment to Deposit Account No. 08-2025. A duplicate copy of this transmittal letter is enclosed.

Respectfully submitted,
Robert Alan Cochran
Olympic Patent Works PLLC


Robert W. Bergstrom
Registration No. 39,906

Enclosures:
Postcards (2)
Transmittal in duplicate
Supplemental Appeal Brief



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent application of:

Inventors: Robert Alan Cochran
Serial No. 09/678,168
Filed: October 02, 2000
For: Method and System for Throttling I/O Request Servicing on Behalf of an I/O Request Generator to Prevent Monopolization of a Storage Device by the I/O Request Generator

Examiner: Akiba K.R. Boyce

Group Art Unit: 3623

Docket No. 10992806-1

Date: May 23, 2005

SUPPLEMENTAL APPEAL BRIEF

Commissioner of Patents and Trademarks
Washington, DC 20231

Sir:

In response to the re-opening of prosecution by the Examiner in the Office Communication mailed February 22, 2005, Applicant hereby requests reinstatement of the Appeal as provided under 37 CFR 1.193(b)(2).

REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

RELATED APPEALS AND INTERFERENCES

Appellant's representative has not identified, and does not know of, any other appeals of interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-20 are pending in the application. Claims 1, 3-8, 10 and 12-17 were rejected and claims 2, 9, 11 and 18-20 were objected to in the Office Action dated February 22, 2005. Appellant's appeal the rejection of claims 1, 3-8, 10 and 12-17 and the objection to claims 2, 9, 11 and 18-20, which are copied in the attached CLAIMS APPENDIX.

STATUS OF AMENDMENTS

No Amendment is enclosed with this brief. The last Response was filed. The last amendment to the claims was filed December 19, 2004.

SUMMARY OF CLAIMED SUBJECT MATTER

Overview

In one embodiment, the present invention is employed by a disk-array controller (Current Application, Figure 2, 206) to prevent monopolization, by remote computers that elect a premium tier of servicing, of I/O request servicing provided by the disk-array controller to remote computers (Current Application, page 10, lines 1-5). In other words, the disk-array controller needs to provide the expected premium-tier servicing of I/O requests to premium-tier remote computers, but needs also to make sure that non-premium-tier computers receive a fair amount of I/O request servicing.

Independent Claim 1

Claim 1 is directed to a method employed by a request servicing device, a disk-array controller in the described embodiment, for fairly servicing electronic requests received from request-generating devices interconnected with the request receiving device. A pricing tier is established for each request-generating device. A maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request

servicing, is established for each request servicing device. The request servicing device maintains an instantaneous rate of request servicing for each request-generating device with a premium pricing tier. After the request servicing device services a request on behalf of a request-generating device, the request servicing device determines a time elapsed during servicing of the request. When the determined time elapsed during servicing of the request is less than the expected time for servicing the request, the request servicing device calculates a remaining time equal to the difference between the expected time and the elapsed time, and then waits for a length of time based on the calculated remaining time prior to servicing another request for the request-generating device.

Dependent Claims 2 – 9

Claim 2 provides details of how the time to wait before servicing another request from a particular request-generating device is computed. Claim 3 is directed to the method of claim 1 wherein the request-generating device is a computer. Claim 4 is directed to the method of claim 1 wherein the request servicing device is a data-storage device. Claim 5 is directed to the method of claim 1 wherein the request servicing device is a disk array data-storage device. Claim 6 is directed to the method of claim 1 wherein a request-generating device specifies the maximum rate of request servicing that it desires. Claim 7 is directed to the method of claim 1 wherein the request servicing device partitions its request-servicing bandwidth among the request-generating devices, and employs this partitioning to establish a maximum rate of request servicing for each request-generating device. Claim 8 is directed to the method of claim 1 wherein the request servicing device dynamically alters the maximum rate of request servicing for a request-generating device. Claim 9 is directed to the method of claim 1 wherein the request servicing device dynamically alters the instantaneous rate of request servicing for a request-generating device after servicing a request in an efficient manner.

Independent Claim 1

Claim 10 is directed to a request servicing device that fairly services electronic requests received from remote request-generating devices. The request servicing device includes a memory that store various parameters associated with each remote request-generating device, including a pricing tier, a maximum rate of request servicing, and an expected time for servicing a request at the maximum rate of request servicing. For each

premium-tier remote request-generating device, the request servicing device continuously computes and maintains, in memory, an instantaneous rate of request servicing. The request servicing device, upon servicing each request, determines the time taken to service the request. When this time is less than the expected time for servicing the request, the request servicing device computes a wait time based on the elapsed time and waits for that wait time before servicing another request for the remote request-generating device.

Dependent Claims 11 – 20

Claim 11 provides details of how the time to wait before servicing another request from a particular request-generating device is computed by the request servicing device. Claim 12 is directed to the method of claim 1 wherein the request-generating device is a computer. Claim 13 is directed to a data-storage-device request servicing device. Claim 14 is directed to a disk array request servicing device. Claim 15 is directed to a request servicing device that allows a request-generating device to specify the maximum rate of request servicing that it desires. Claim 16 is directed to a request servicing device that partitions its request-servicing bandwidth among the request-generating devices, and employs this partitioning to establish a maximum rate of request servicing for each request-generating device. Claim 17 is directed to a request-servicing device that dynamically alters the maximum rate of request servicing for a request-generating device. Claims 18-19 are directed to a request-servicing device that dynamically alters the instantaneous rate of request servicing for a request-generating device after servicing a request in an efficient manner. Claim 20 is directed to a request-servicing device that computes the length of time to wait before servicing another request for a particular request-generating device by an efficient method.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether claim 1 is indefinite under 35 U.S.C. § 112, second paragraph, for failing to define the term "pricing tier."
2. Whether claims 1, 3, 4-6, 10, and 12-15 are patentable over Freeland et al., U.S. Patent No. 4,262,331 in view of Kilkki et al., U.S. Patent No. 6,011,778.

3. Whether claims 7, 8, 16, and 17 are patentable over Freeland et al., U.S. Patent No. 4,262,331 in view of Kilkki et al., U.S. Patent No. 6,011,778 and further in view of Storch et al., U.S. Patent 5,920,846.

ARGUMENT

Claims 1-20 are currently pending in the current application. In the Office Action dated February 22, 2005 ("Office Action"), the Examiner re-opened prosecution and rejected claim 1 under 35 U.S.C. § 112, second paragraph, for failing to define the term "pricing tier," rejected claims 1, 3, 4-6, 10, and 12-15 under 35 U.S.C. § 103(a) as being unpatentable over Freeland et al., U.S. Patent No. 4,262,331 ("Freeland") in view of Kilkki et al., U.S. Patent No. 6,011,778 ("Kilkki"), and rejected claims 7, 8, and 16-17 under 35 U.S.C. § 103(a) as being unpatentable over Freeland in view of Kilkki in further view of Storch et al., U.S. Patent No. 5,920,846 ("Storch"). Appellant's representative respectfully traverses these rejections.

ISSUE 1

1. Whether claim 1 is indefinite under 35 U.S.C. § 112, second paragraph, for failing to define the term "pricing tier."

In the Office Action, the Examiner states:

In this case, the term "pricing tier" in claim 1 is a relative term that renders the claims indefinite. The term 'pricing tier' is not defined by the claim, the specification does not provide a standard for ascertaining the requisite degree, and one of ordinary skill in the art would not be reasonably apprised of the scope of the invention. In addition, once this "pricing tier" is established, it is not utilized in any way to process service requests. This pricing tier" is therefore irrelevant.

Appellant's representative could not more strongly disagree with this statement. First, the term "pricing tier" is used 4 times in the current application, and is quite well defined. For example, on page 16, the enumeration type "price," which occurs in the first line of the C++-like pseudocode on page 14, is described as follows:

In the embodiment of the present invention described in the C++-like pseudocode, the remote computer may contract either for **basic service**, essentially representing the simple throttling mechanism illustrated in Figure 5 in which the remote computer that

unevenly generates I/O requests may receive substantially less than the maximum rate of I/O request servicing, or a **premium service** in which the remote computer receives a rate of I/O request servicing as close to the maximum contracted-for rate possible using the above-described sliding window throttling mechanism. Thus, the enumeration "**price**" on line 1, above, contains values representing the **two pricing tiers** available to remote computers, namely "**BASIC**" representing I/O request servicing according to the simplified throttling mechanism, and "**PREMIUM**" representing I/O request servicing according to the sliding window throttling mechanism. (emphasis added)

The term "pricing tier" occurs again in the second-to-last line at the bottom of page 18, on the last line of page 23, on the fourth line of page 24, and frequently occurs in the claims. As stated in the above quoted portion of the Specification, a pricing tier represents a tier of service, and two tiers of service, or pricing tiers, are employed in the described embodiment: a basic pricing tier; and a premium pricing tier. Not only is the term "pricing tier" defined in the current application, C++-pseudocode is included that features an explicitly enumeration of pricing tier, *price*, used repeatedly in the C++-pseudocode. Moreover, the term "pricing tier" is claimed to be utilized for processing service requests in claim 1, as can be readily seen in the following portion of claim 1, with emphasis added:

establishing a pricing tier for each request generating device, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing;
for each request generating device with a premium pricing tier,
maintaining an instantaneous rate of request servicing by the request-servicing device;

There is no basis in the current application or the current claims for the Examiner's statement. Furthermore, the Examiner's statement reveals a common misconception – that the *claims* need to define terms used in the claims. This is not the case. Terms that need definitions are generally defined in the Specification. Otherwise, claims would be encyclopedic and impossible to state or read. It would be an unfair burden to Appellant, in time and expense, for prosecution to be re-opened based on this new, unfounded rejection.

ISSUE 2

2. Whether claims 1, 3, 4-6, 10, and 12-15 are patentable over Freeland et al., U.S. Patent No. 4,262,331 in view of Kilkki et al., U.S. Patent No. 6,011,778.

The patentability of the currently claimed invention over Kilkki has been discussed, in detail, in the originally filed Appeal Brief. Portions of that discussion are repeated, below, following a brief overview of the Currently Claimed Invention. The originally filed Appeal Brief may be consulted for additional details and citations. Following this repeated discussion of Kilkki, the currently claimed invention is compared to the system and method disclosed by Freeland.

Currently Claimed Invention

Claim 1 is directed to a method for servicing electronic requests employed by a request-servicing device. In an embodiment described in the current application in detail, the request-servicing device is a disk-array controller running within a disk array. The disk-array controller receives I/O requests from remote computers and services the received requests, returning data and/or status related to the I/O requests to the requesting, remote computers. As clearly claimed in claim 1, the request-servicing device, or disk-array controller in the described embodiment, establishes and maintains, for each remote computer on behalf of which the request-servicing device services I/O requests, a pricing tier, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing. After the request-servicing device services a request on behalf of a remote, premium tier computer, the disk-array controller computes an elapsed time for servicing of the request. When the computed elapsed time is less than the expected time for servicing the I/O request, the request-servicing device computes a difference between the elapsed time and the expected time and waits for a time equal to the computed difference before servicing another I/O request on behalf of the remote computer. This allows a premium tier computer to receive a greater rate of request servicing than contracted for, while basic tier computers obtain only up to a rate of request servicing that they contracted for. However, premium tier computers are not allowed to monopolize the request-servicing device, because the request servicing device always provides for processing intervals for basic tier computers when there are basic-tier-computer requests to service.

Kilkki

Kilkki, in contrast to the currently claimed invention, is directed to a timer-based method employed by an access node for measuring transmission rates of information by a network and for adjusting message priority assignments by an access node to achieve a

nominal bit rate for transmission by the network. Kilkki's method is best understood with reference to Kilkki's Figure 2. Kilkki's disclosed system involves a user (Figure 2, 20) that transmits messages, or, in Kilkki's terminology, transmits cells, through a user/network interface ("UNI") (Figure 2, 24) to a first node (Figure 2, 32) in a computer network (Kilkki, column 7, lines 16-18) that, in turn, transmits the cells to additional computers (Figure 2, 34) in the network; a final network node 34 forwarding the cells to destination computers (Figure 2, 36). The user interacts, through the UNI 24, with a network operator (Figure 2, 22), presumably another computer, to negotiate (or, presumably, to purchase) a nominal bit rate ("NBR") for transmitting data through the network (Kilkki, column 7, lines 45-46). Kilkki does not appear to explicitly state whether the UNI is a software program running on the user's computer, or is instead a separate access node distinct from the user's computer.

Taking the user/UNI combination as the request-generating device, and the first network node (Figure 2, 32) as the request-servicing device, since the request-generating device, or user/UNI, is disclosed in Kilkki as generating cell-transmission requests and forwarding those requests to the first network node, the nominal-bit-rate method of Kilkki can be compared to Appellant's currently claimed invention. In Kilkki's system, the UNI attempts to achieve the nominal bit rate ("NBR"), negotiated for with the network operator by the user, for transmitting data through the network. The NBR represents an expected, but not a guaranteed, bit rate associated with the user (Kilkki, column 3, lines 32 – 34). The UNI 24 sets a timer to an expected time upon transmitting a first cell to the first network node 32. The timer value is evaluated by the UNI at the point in time that the UNI transmits a second cell to the first network node. The evaluation of the timer value provides a difference between the expected time interval between transmission of the first and second cells and the actual time interval between transmission of the first and second cells. In other words, the elapsed time calculated *by the request-generating device* is an elapsed time between issuing electronic cell-transmission requests by the *request-generating device*. The UNI assigns a priority level to the second message based on this difference (Kilkki, column 3, lines 53-65; column 6, lines 11-35). The first network node 32 receives the cell from the UNI, and decides, based on the priority assigned to the cell by the UNI, whether or not to transmit the cell through the network, or to instead discard the cell:

Claim 1 of the current application, by contrast, claims a method for fairly servicing, *by a request-servicing device*, electronic requests received by the request-servicing device from request-generating devices interconnected with the request receiving device. In

the first step of the claimed method, the *request-servicing device* (1) establishes a pricing tier for each request-generating device; (2) establishes a maximum rate of request servicing; and (3) establishes an expected time for serving a request at the maximum rate of request servicing. By contrast, in Kilkki's disclosed system, the network operator (Figure 2, 22), and not the request-servicing device (first network node 32), establishes the NBR, or, in other words, the maximum rate of request servicing. In Kilkki's system, the *request-generating device* (user 20 and UNI 24) computes an *expected time between issuing cell-transmission requests* – the only expected time mentioned in Kilkki. Note that this is not an expected time for servicing a particular request, or for servicing requests in general, as in the currently claimed method. The request-servicing device, in the currently claimed method, establishes a pricing tier, a maximum rate of request servicing, and an expected time for servicing a request. In Kilkki's system, by contrast, neither a request-generating device nor a request-servicing device establishes the NBR. In Kilkki's system, the request-generating device, rather than the request-servicing device, establishes an expected time. However, the expected time is an expected time interval between making cell-transmission requests of the first network node by the user/UNI, and is not an expected time for servicing a request. Kilkki explicitly states, in several places, that the network node, or request-servicing device, needs to know very little about the user/UNI devices: In other words, the first network node 32 knows almost nothing about user/UNIs. In particular, the request-servicing device in Kilkki does not compute or maintain information with respect to expected service times or service rates. As discussed above, any such activities are the responsibility of the *request-generating device* in Kilkki's system.

The second step of current claim 1 states that the *request-servicing device* maintains an instantaneous rate of request servicing by the request-servicing device for each request-generating device with a premium pricing tier. In the described embodiment of the current invention, the request-servicing device computes and maintains, based on a sliding-window history of request servicing, a recent request servicing rate with respect to each remote client. That is quite different from requesting that a user supply a desired instantaneous NBR to the network. The request-servicing device in Kilkki's system maintains almost no information about remote user/UNIs. There is no indication that even the request-generating device in Kilkki's system, a user/UNI, computes and maintains an instantaneous rate of request servicing in Kilkki's system. At most, in Kilkki's system, the user/UNI computes a difference between an interval from sending a first request and sending

a second request and an expected interval between sending the first and second requests.

In the final steps of claim 1, when the time elapsed during servicing of a request is less than the expected time for serving a request established for the request-generating device *by the request-servicing device*, the request-servicing device, or disk-array controller in the described embodiment of the current invention, calculates a remaining time equal to the difference between the expected time for serving a request established for the request-generating device and the time elapsed during servicing of the request, and *waits for a length of time based on the calculated remaining time* prior to servicing another request for the request-generating device. In Kilkki's system, neither the request-generating device (user/UNI) nor the request-servicing device (network node) waits for any length of time. In Kilkki's system, the request-generating device computes a priority to associate with the next cell-transmission request based on the time interval between sending that request and sending the previous request. The request-generating device then sends the priority-associated cell-transmission request. It does not wait for a computed period of time before doing so. In Kilkki's system, the network node determines, based on the priority associated with the cell-transmission request, whether to transmit the cell or discard the cell. The request-servicing device does not wait for a computed length of time. The request-servicing device in Kilkki may, based on the priority assigned to a cell, not service the request at all.

The claimed invention involves a method, and a system embodying the method, by which a request-servicing device prevents monopolization of request servicing bandwidth provided by the request-servicing device by a remote request-generating device, as clearly stated even in the title of the current application. By contrast, Kilkki discloses a method by which a request-generating device attempts to ensure that the request-generating device obtains adequate request servicing from a request-servicing, multiple-device system.

Freeland

Freeland, like Kilkki, involves *request-generating-device* control of a request-servicing device, both unlike the currently claimed invention, in which a *request-servicing device* monitors and adjusts the rate at which it services requests from request-generating devices. Freeland is directed to a computer system comprising a CPU and a number of peripheral processing units ("PEs"). The PEs compete with one another for servicing, by the CPU, of various processing tasks (Freeland, column 1, lines 17-23). In Freeland's disclosed system, registers are included in each PE for storing the times of commencement of the last

service period and the current service period by the CPU on behalf of the PE (Freeland, column 2, lines 1-5). The PE can compute the interval between commencements of the previous and current service times, and compare the computed interval to an expected interval (Freeland, column 1, lines 5-9). When the computed interval falls below the expected interval, the PE sets the amount of processing next requested from the CPU to some desired number (Freeland, column 2, lines 13-16), and when the expected interval falls above the expected interval, the PE adjusts the amount of processing next requested from the CPU downward (Freeland, column 2, lines 16-21). This is also described beginning on line 63 of column 4 to line 3 of column 5, where the number of processing steps requested is either a maximum number N_m or an expected number N_E less than N_m . Like Kilkki, Freeland does not teach, mention, or suggest increasing the request of services above the maximum amount for a period of time in order to achieve an expected rate. Freeland lacks a means for doing so, because Freeland's PEs only have registers for computing a single interval, and comparing that to a single expected value. Freeland cannot monitor a rate of processing over multiple processing sessions. Freeland does not teach, mention, or suggest the sliding window approach of the current invention, or any approach, for computing the claimed an instantaneous rate of request servicing. In the currently claimed invention, the request-servicing device establishes a pricing tier to assign to each remote computer, but, in Freeland, the request-generating device is apparently manufactured or configured to have a priority, which the request-generating device uses to increase the number of requested operations from the CPU by multiplication (Freeland, column 2, lines 21-26). In the currently claimed invention, a time elapsed during servicing of requests is computed by a request-generating device, while each PE, or request generating device, in Freeland computes an interval between when the previous servicing of requests on behalf of the PE by the CPU commenced and when current servicing of requests on behalf of the PE by the CPU commenced. In the currently claimed invention, the request-servicing device waits for a time computed based on the computed instantaneous rate of request servicing for a particular remote computer before beginning to service another request, while, in Freeland's system, the request-generating device adjusts the amount of servicing that the request-generating device will request next time.

Claim 1 is provided below, for comparison to Freeland's disclosed system:

1. A method for fairly servicing, by a request servicing device, electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the method comprising:

establishing a pricing tier for each request generating device, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing;

for each request generating device with a premium pricing tier, maintaining an instantaneous rate of request servicing by the request servicing device;

following servicing of each request from a request generating device by the request servicing device,

determining a time elapsed during servicing of the request;

when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device,

calculating a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request; and

waiting for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

Considering each element of claim 1 of the current application, it is readily apparent that Freeland neither teaches nor suggests a single element of claim 1. Freeland does not establish a pricing tier, a maximum rate of request servicing, or an expected time for request servicing, and Freeland's request servicing device, the CPU, establishes nothing at all with respect to servicing of requests on behalf of any PE. Instead, Freeland's PEs have priorities built into them or configured into them, and the priorities are not based on a contracted-for servicing rate associated with a pricing tier. Freeland's CPU does not establish an expected time for servicing a request, but, instead, Freeland's PEs establish an expected time between commencements of two request servicing periods. This is significantly different, because a rate of request servicing directly reflects the amount of work performed by the request-servicing device for a particular remote computer in the currently claimed system, while the interval between commencement of servicing requests by a request-servicing device in Freeland's system does not, since the amount of processing requested by a PE during each processing period can vary. In Freeland's system, each PE includes the special registers for computing intervals, and computes an amount of processing to request in the same fashion, while, in the currently claimed invention, the request-servicing device maintains an instantaneous rate of request servicing for premium-pricing-tier devices. Neither Freeland's CPU nor Freeland's PEs maintain an instantaneous rate of request servicing. Neither Freeland's CPU nor Freeland's PEs determine a time elapsed during servicing of each request

from a request-generating device by the request-servicing device. Neither Freeland's CPU nor Freeland's PEs calculate a remaining time equal to the difference between an expected time for serving a request established for the request-generating device and the time elapsed during servicing of the request. Finally, Freeland's CPU does not wait for a length of time based on the calculated remaining time prior to service another request for the request-generating device following servicing of each request from a request-generating device by the request-servicing device. Freeland's CPU performs requested services by PEs on demand.

In Appellant's representative's opinion, a combination of Freeland and Kilkki does not come any closer to the currently claimed invention than Kilkki alone. The Examiner's new 35 U.S.C. § 103(a) rejection based on a combination of Freeland and Kilkki is unfounded, and it would be an unfair burden to Appellant, in time and expense, for prosecution to be re-opened based on this new rejection.

ISSUE 3

3. Whether claims 7, 8, 16, and 17 are patentable over Freeland et al., U.S. Patent No. 4,262,331 in view of Kilkki et al., U.S. Patent No. 6,011,778 and further in view of Storch et al., U.S. Patent 5,920,846.


Neither Freeland, nor Kilkki, nor Freeland and Kilkki in combination teach or suggest the currently claimed invention, as discussed in the previous section directed to Issue 2. The reference Storch was discussed in the originally filed Appeal Brief. Storch is unrelated to the currently claimed invention. In the Office Action, the Examiner has simply repeated the original citation of Storch, without responding to Appellant's representative's arguments. The Examiner has not introduced a substantial new argument or a substantially more related reference, and therefore, in Appellant's representative's respectfully offered opinion, it would be an unfair burden to Appellant, in time and expense, for prosecution to be re-opened based on this new rejection.

CONCLUSION

The Examiner's new 35 U.S.C. § 112, second paragraph, rejection of claim 1 is unfounded. The newly cited reference, Freeland, is essentially unrelated to the currently claimed invention, and does not, alone or in combination with previously cited references, advance the Examiner's rejection of claims 1, 3-8, 10, and 12-17. Appellant appreciates the Examiner's conditional allowance of claims 2, 9, 11, and 18-20, but continues to maintain patentability of the currently claimed invention over any combination of the cited references. For this reason, Appellant respectfully requests that the Appeal of the Examiner's rejections in the Office Action mailed on April 05, 2004, finally rejecting claims 1-20, be maintained.

Applicant respectfully submits that all statutory requirements are met and that the present application is allowable over all the references of record. Therefore, Applicant respectfully requests that the present application be passed to issue.

Respectfully submitted,
Robert Alan Cochran
OLYMPIC PATENT WORKS PLLC

By 
Robert W. Bergstrom
Reg. No. 39,906

Olympic Patent Works PLLC
P.O. Box 4277
Seattle, WA 98104
206.621.1933 telephone
206.621.5302 fax

CLAIMS APPENDIX

1. A method for fairly servicing, by a request servicing device, electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the method comprising:

establishing a pricing tier for each request generating device, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing;

for each request generating device with a premium pricing tier, maintaining an instantaneous rate of request servicing by the request servicing device;

following servicing of each request from a request generating device by the request servicing device,

determining a time elapsed during servicing of the request;

when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device,

calculating a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request; and

waiting for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

2. The method of claim 1 wherein the length of time based on the calculated remaining time is determined to be:

the calculated remaining time for a request generating device for which the established pricing tier is a basic pricing tier;

the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is equal to the maximum rate of request servicing established for the request generating device;

greater than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is greater than the maximum rate of request servicing established for the request generating device; and

less than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is less than the maximum rate of request servicing established for the request generating device.

3. The method of claim 2 wherein the request generating device is a computer.

4. The method of claim 2 wherein the request servicing device is an electronic data storage device.

5. The method of claim 4 wherein the electronic data storage device is a disk array.

6. The method of claim 1 wherein the maximum rate of request servicing is established via specification of a maximum rate of request servicing by the request generating device.

7. The method of claim 1 wherein the maximum rate of request servicing is established by partitioning the capacity of the request servicing device among the request generating devices in order to provide, when possible, each request generating device with a maximum rate of request servicing specified by the request generating device, and otherwise to provide each request generating device with a maximum rate of request servicing proportional to a maximum rate of request servicing specified by the request generating device.

8. The method of claim 1 wherein the request servicing device may dynamically alter the maximum rate of request servicing provided to one or more request generating devices in accordance with a rate at which the request servicing device receives requests and according to the request servicing capacity of the request serving device.

9. The method of claim 1 wherein maintaining an instantaneous rate of request servicing by the request servicing device further comprises:

initially setting the instantaneous rate of request servicing for a request generating device to one request divided by the expected time for serving a request at the maximum rate of request servicing established for the request generating device;

increasing the instantaneous rate of request servicing for the request generating device by one following servicing of a request generated by the request generating device; and

decreasing the instantaneous rate of request servicing for the request generating device by one at regular intervals of time.

10. A request servicing device that fairly services electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the request servicing device including:

a memory that contains an established maximum rate of request servicing, an expected time for serving a request at the maximum rate of request servicing, and a pricing tier for each request generating device and that contains, for each request generating device with a premium pricing tier, an instantaneous rate of request servicing by the request servicing device; and

control functionality that services electronic requests received from the request generating devices and that, following servicing of each request from a request generating device by the request servicing device, determines a time elapsed during servicing of the request so that, when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device, the control functionality calculates a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request and waits for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

11. The request servicing device of claim 1 wherein the length of time based on the calculated remaining time is determined by the request servicing device to be:

the calculated remaining time for a request generating device for which the established pricing tier is a basic pricing tier;

the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request

servicing is equal to the maximum rate of request servicing established for the request generating device;

greater than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is greater than the maximum rate of request servicing established for the request generating device; and

less than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is less than the maximum rate of request servicing established for the request generating device.

12. The request servicing device of claim 10 wherein the request generating device is a computer.

13. The request servicing device of claim 10 wherein the request servicing device is an electronic data storage device.

14. The request servicing device of claim 10 wherein the electronic data storage device is a disk array.

15. The request servicing device of claim 10 wherein the maximum rate of request servicing is established via specification of a maximum rate of request servicing by the request generating device.

16. The request servicing device of claim 10 wherein the maximum rate of request servicing is established by partitioning the capacity of the request servicing device among the request generating devices in order to provide, when possible, each request generating device with a maximum rate of request servicing specified by the request generating device, and otherwise to provide each request generating device with a maximum rate of request servicing proportional to a maximum rate of request servicing specified by the request generating device.

17. The request servicing device of claim 10 wherein the request servicing device may dynamically alter the maximum rate of request servicing provided to one or more request generating devices in accordance with a rate at which the request servicing device receives requests and according to the request servicing capacity of the request serving device.

18. The request servicing device of claim 10 wherein maintaining an instantaneous rate of request servicing by the request servicing device further comprises:

initially setting the instantaneous rate of request servicing for a request generating device to one request divided by the expected time for serving a request at the maximum rate of request servicing established for the request generating device;

increasing the instantaneous rate of request servicing for the request generating device by one following servicing of a request generated by the request generating device; and

decreasing the instantaneous rate of request servicing for the request generating device by one at regular intervals of time.

19. The request servicing device of claim 18 wherein separate execution threads in a firmware or software implementation of control functionality within the request servicing device increase the instantaneous rate of request servicing and decrease the instantaneous rate of request servicing.

20. The request servicing device of claim 10 wherein the length of time based on the calculated remaining time determined by the request servicing device to be greater than the calculated remaining time for a request generating device is further determined to be a ratio multiplied by the calculated remaining time, the ratio calculated by dividing the instantaneous rate of request servicing by the expected time for serving a request, and wherein the length of time based on the calculated remaining time determined by the request servicing device to be less than the calculated remaining time for a request generating device is further determined to be a ratio multiplied by the calculated remaining time, the ratio calculated by dividing one by the difference between the expected time for serving a request and the instantaneous rate of request servicing.